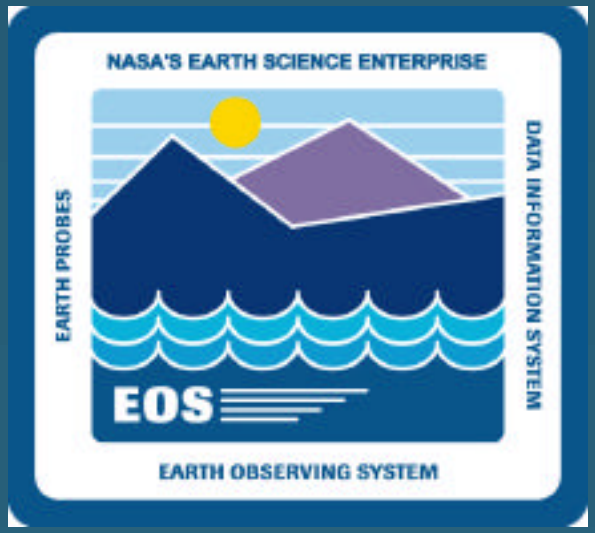

Multi-Host Support for Media Access

Don Brown

Raytheon

Requirements Summary

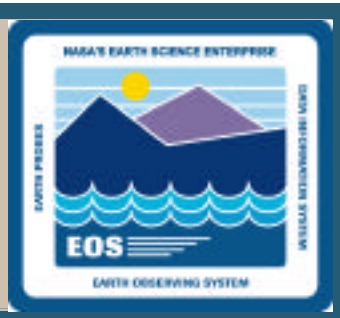

- Allow requests for a single media type to be handled by devices on different hosts.
- Allow for the re-routing of a request that has encountered a device error
- Allow operations staff to view the progress of ftp transfers
- Provide the operations staff with more status on media transfers
- Ensure that media resources are freed up on client termination
- Ensure that resources are made available on server start up
- Allow DAACs to handle peak loads

Key Design Drivers

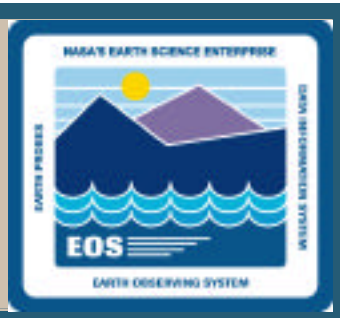

- Maintain current client interfaces
- Use the STMGT Request Manager architecture
- Performance Requirements
  - Meet peak load needs
  - Meet data availability requirements
  - Maximize throughput to media device
- Provide scalable media resource access
- Monitor the state of the STMGT servers
- Restructure Media Configuration GUIs based on DAAC feedback

New SW Components

- One new CI - STMGT Request Manager
  - STMGT Request Manager Processes
  - One STMGT Monitor Process for each Resource Manager
  - Additional Resource Manager(s) per media type

New HW Components

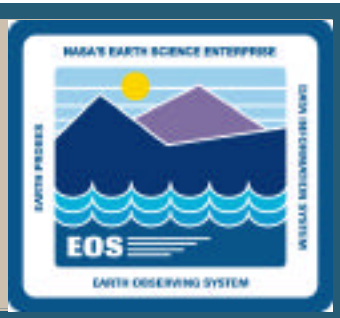

- No new hardware
- RAID attached to DIP hardware will be nfs mounted to facilitate request re-routing
- Multiple hosts with media devices attached

Interaction Diagram

Ftp Distribution Example

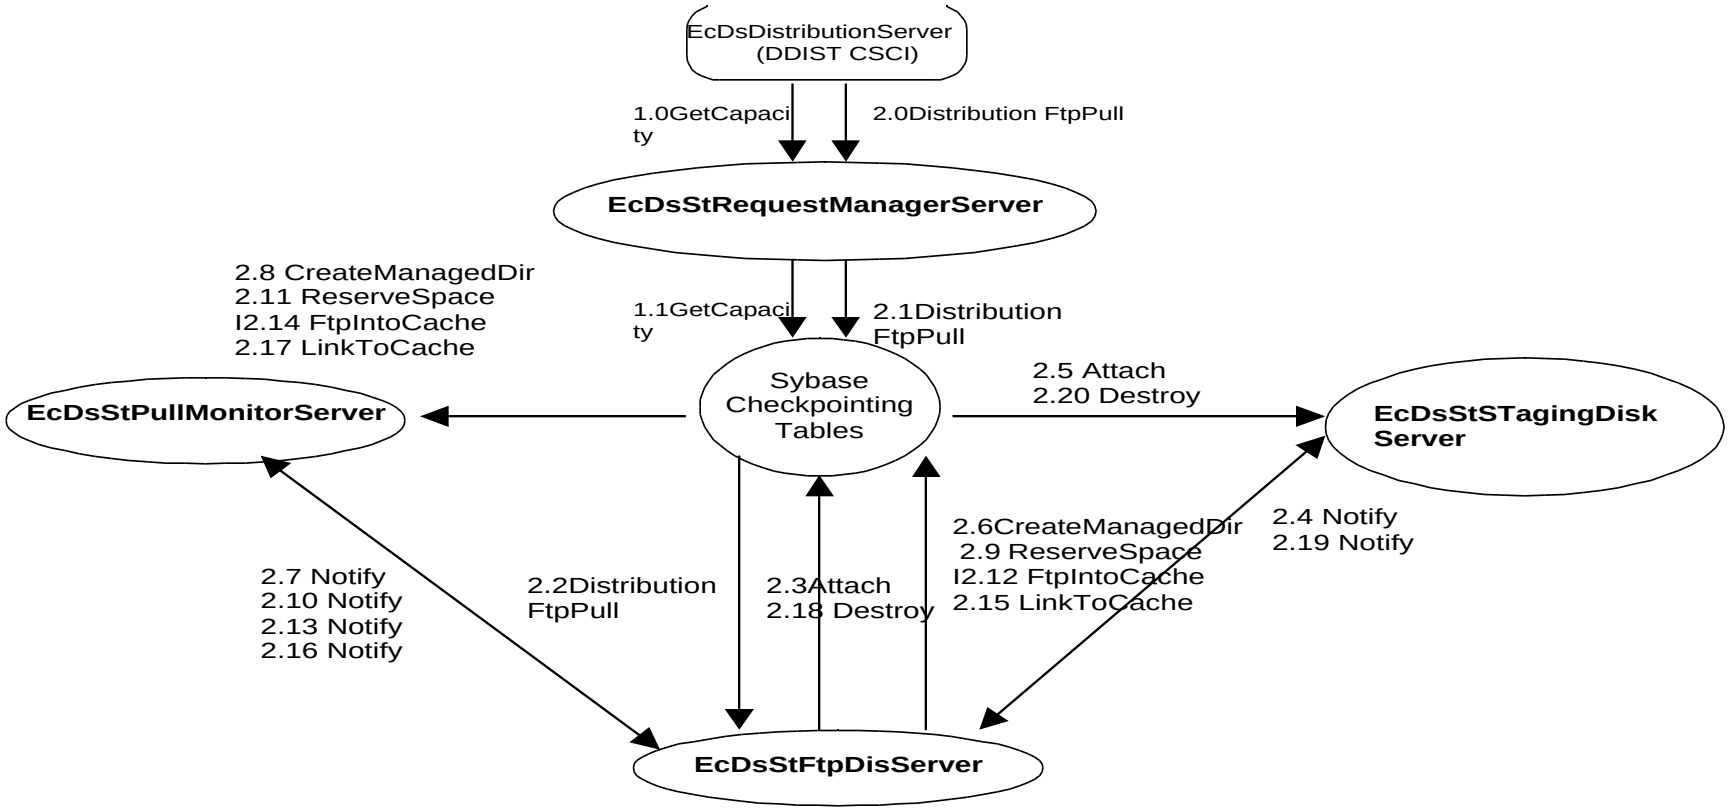

Reference DID 305 Figure 4.1.3.4-1

End User Interaction

STMGT Configuration Tab

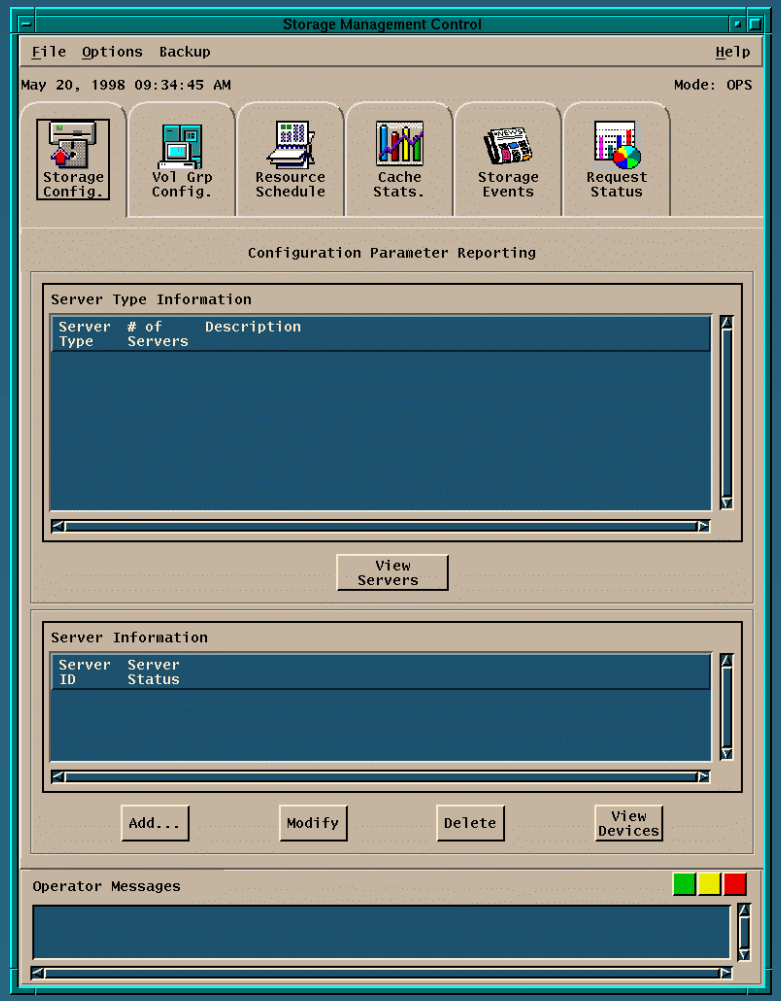

DID 609 4.10.2-1

End User Interaction

Configure Stacker Server

DID 609 4.10.2-3

End User Interaction

Stacker Configuration

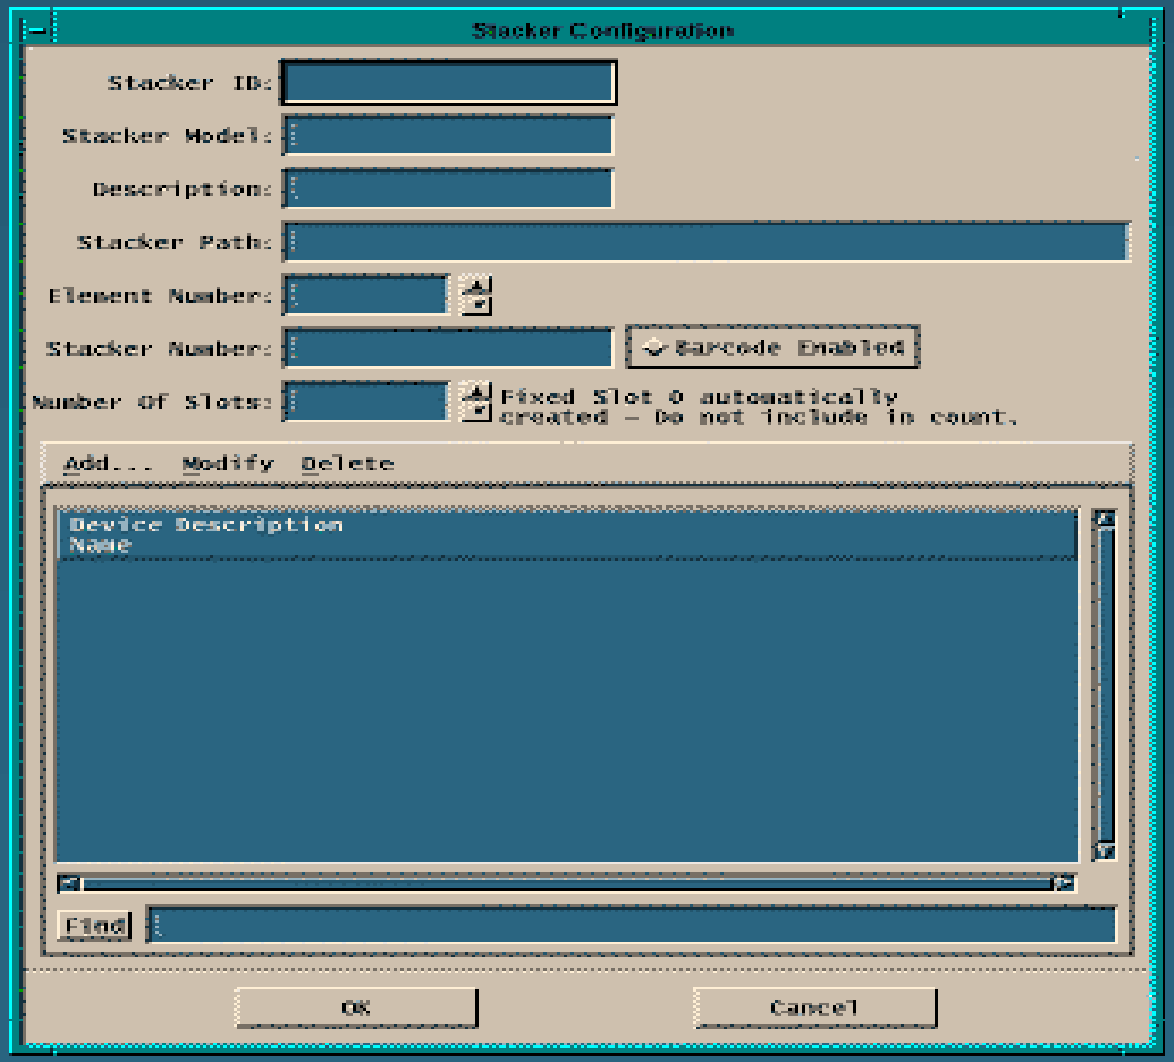

DID 609 4.10.2-3.1

End User Interaction

Device Drive Configuration

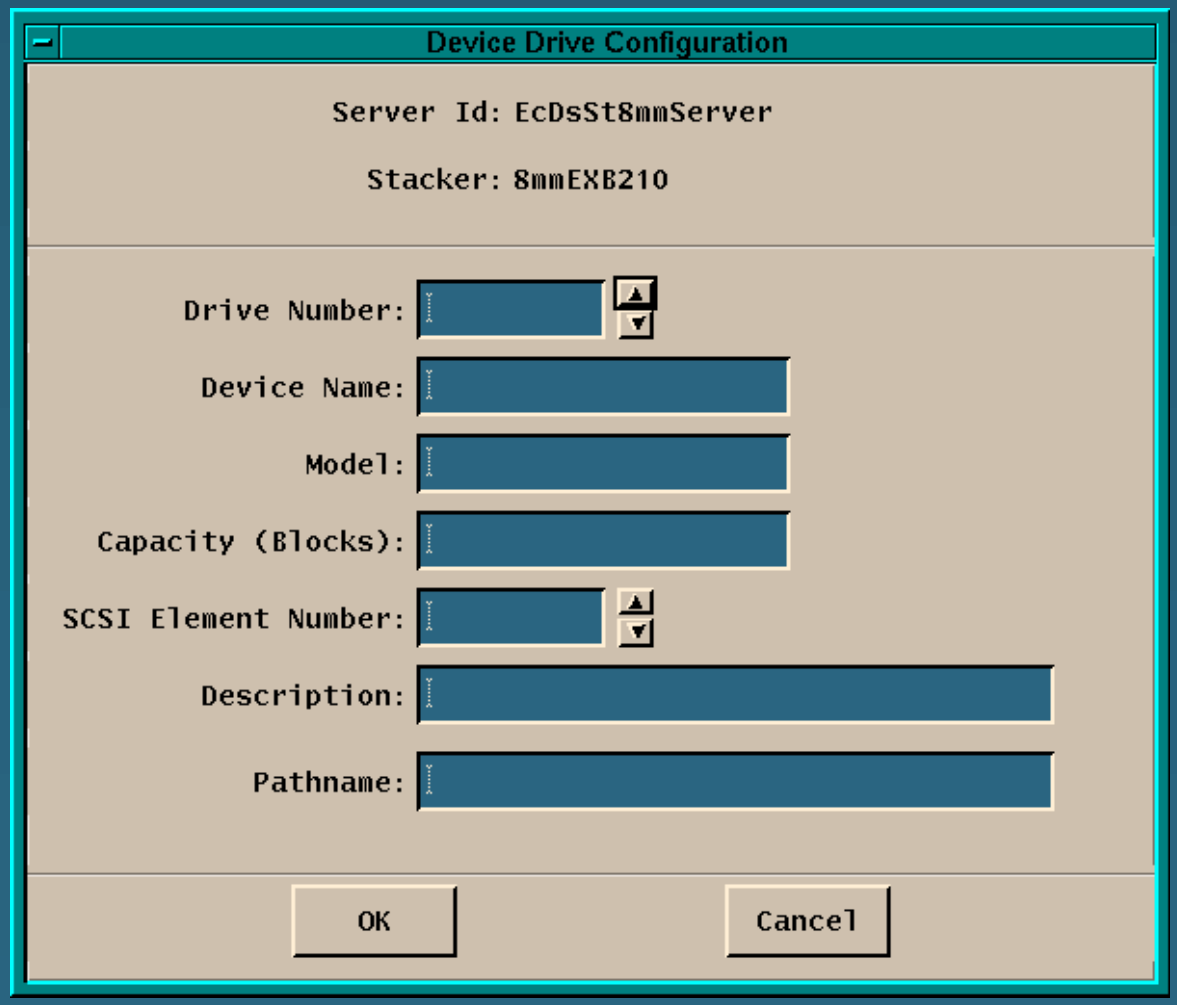

DID 609 4.10.2-3.2

End User Interaction

Resource Scheduling Tab

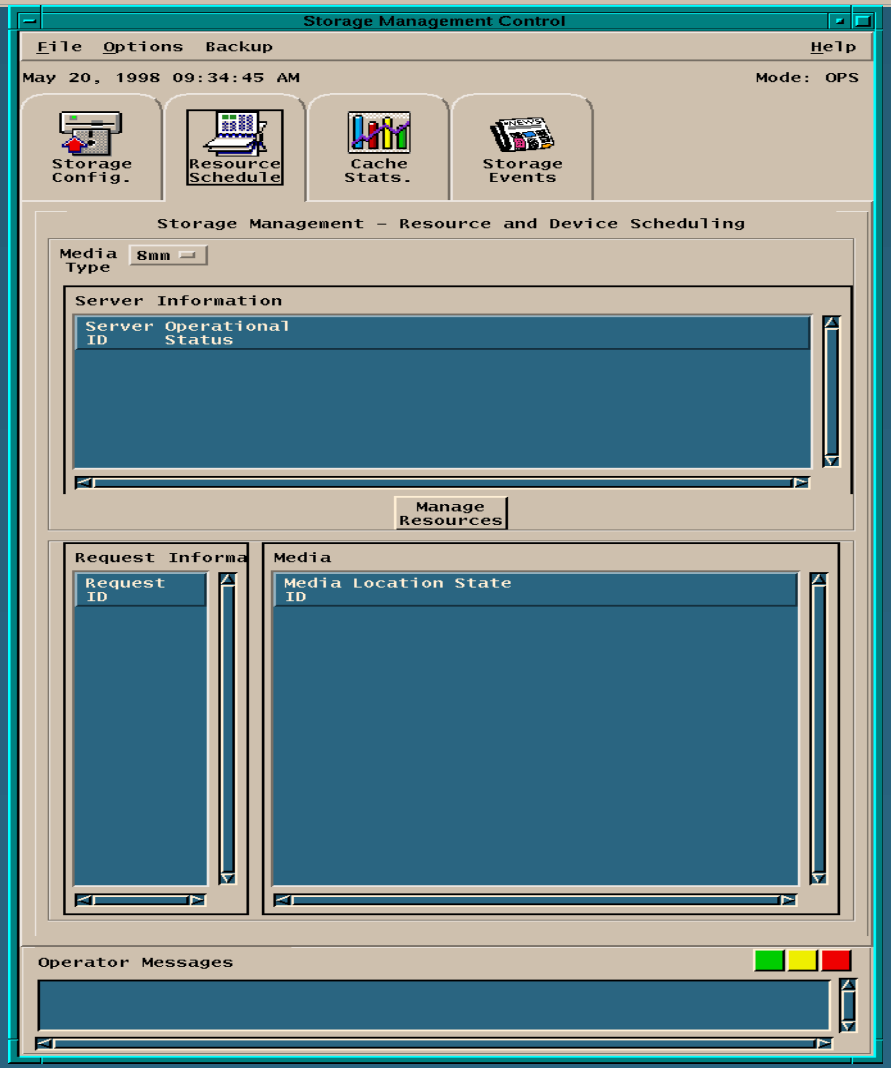

DID 609 4.10.2-14

End User Interaction

Stacker Status Pop up

DID 609 4.10.2-15

End User Interaction

Tape Information

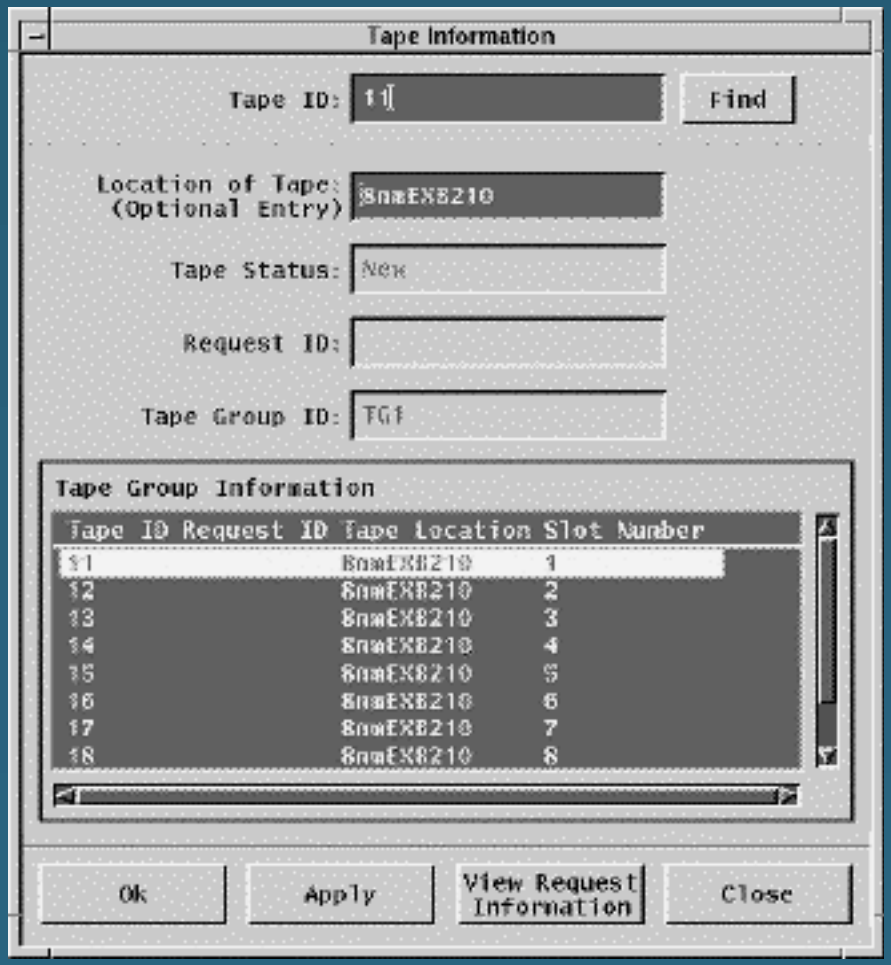

DID 609 4.10.2-22

DAAC Operational Impacts

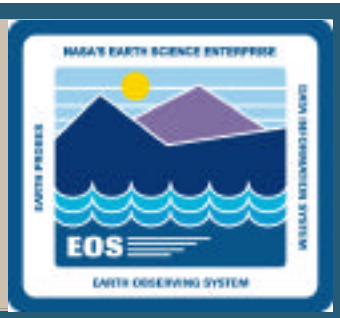

- Multiple servers for a given media type may have to be configured
- Configuration of Tape Stackers change based on DAAC feedback
- Progress of Ftp Transfers will be observable
- Media Stackers or devices for one media type will have to be apportioned to different hosts.

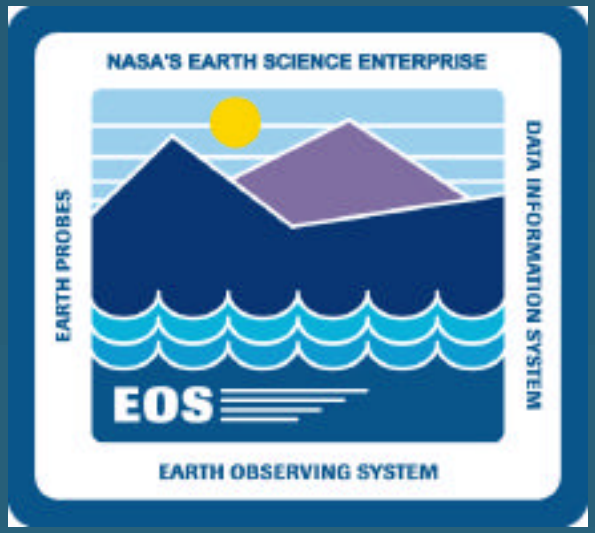

Granule Deletion Administration

Adrienne Dupree

Raytheon

Requirements Summary

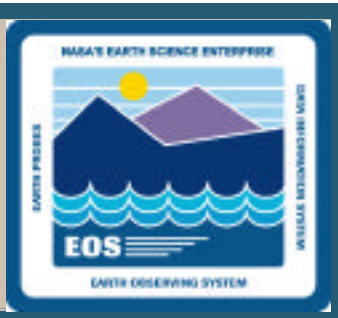

- Granule Deletion Administration supported by requirements in SDSRV and DMS
  - Delete granules from inventory and archive or just archive (DFA)
  - Specify criteria to delete granules
    - -ESDT ShortName, ESDT Version ID, Granule Insert Time Range
    - -File containing ESDT ShortName, ESDT Version ID, Local Granule ID
    - -File containing SDSRV Granule Ids (GEOIDS)
  - Deletion of associated granules if not referenced by any other granule
    - -Browse, PH, QA
  - Capability to display list of granules that will be deleted prior to deletion

Key Design Drivers

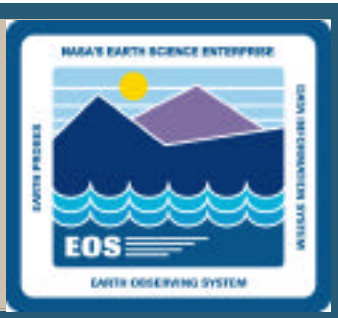

- Allow operators to delete products on demand using a command line utility
- Enforce referential integrity constraints during granule deletion
- Gracefully handle errors caused by concurrent deletions, searches, and acquires
- Allow for retry of an interrupted granule deletion operation

New SW Components

- New standalone program EcDsGranuleDelete

Interaction Diagram

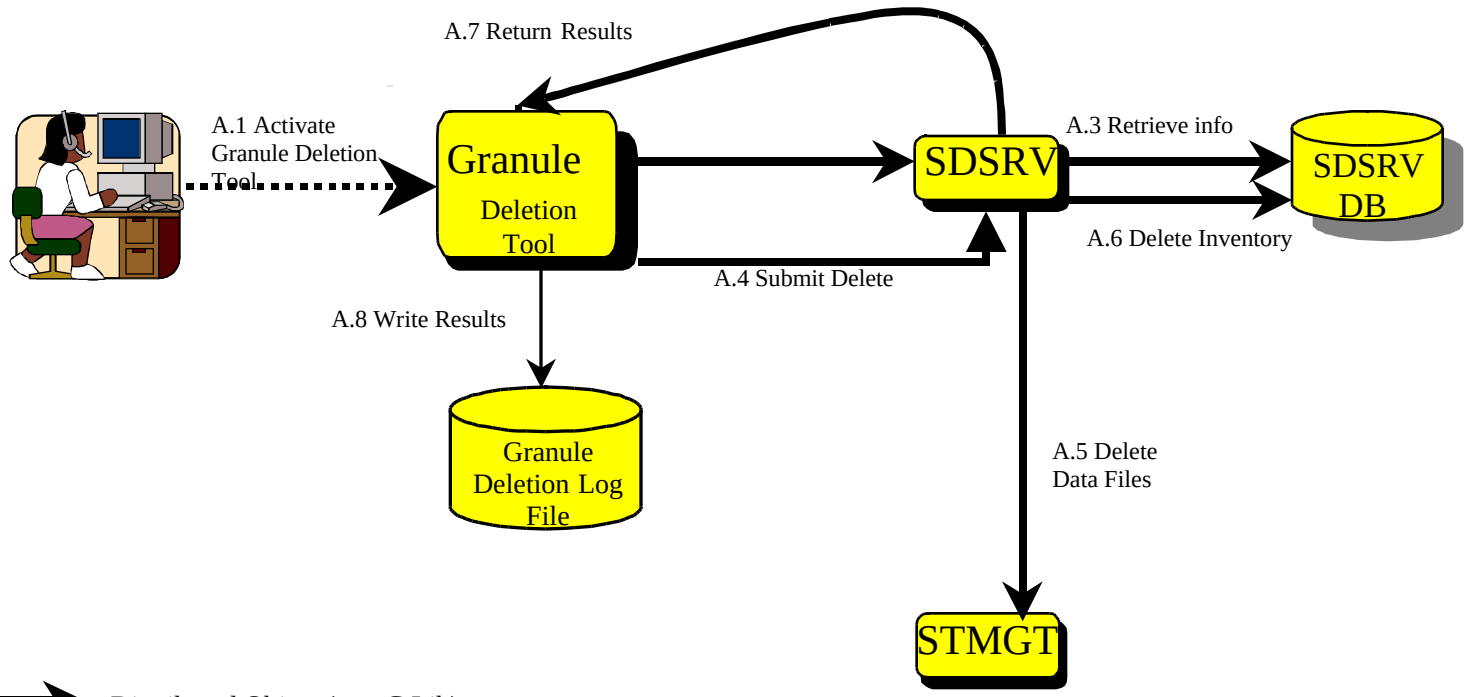

Physical Delete Without Confirmation

Interaction Diagram

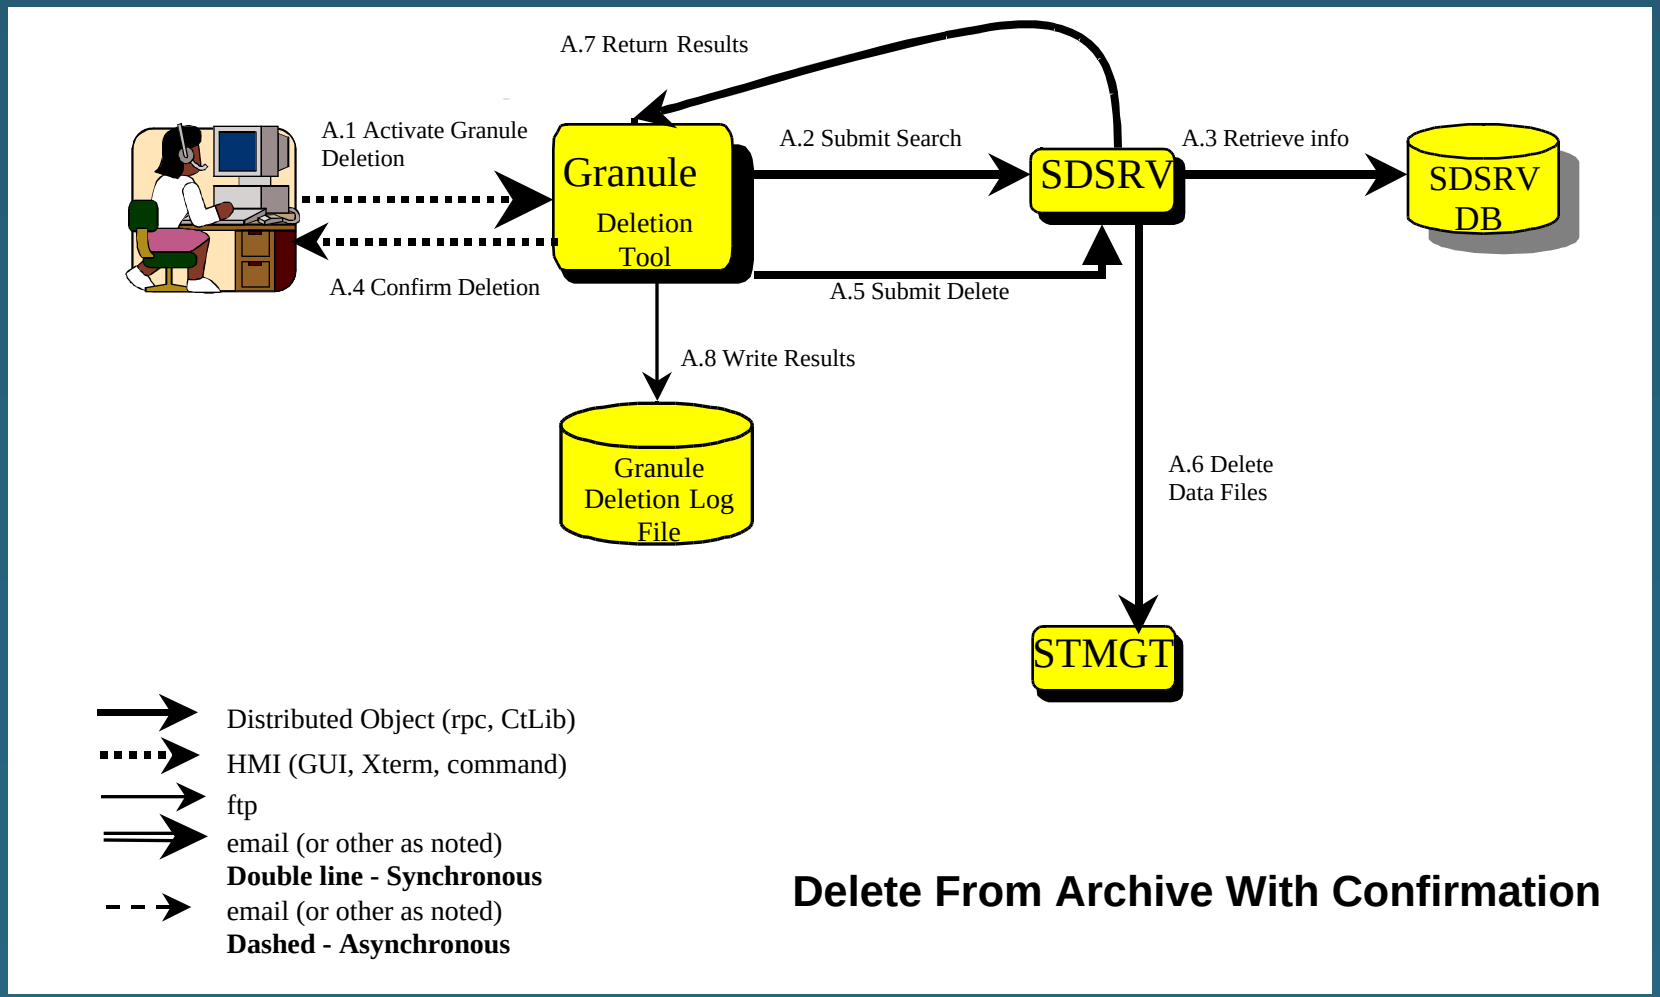

Interaction Diagram

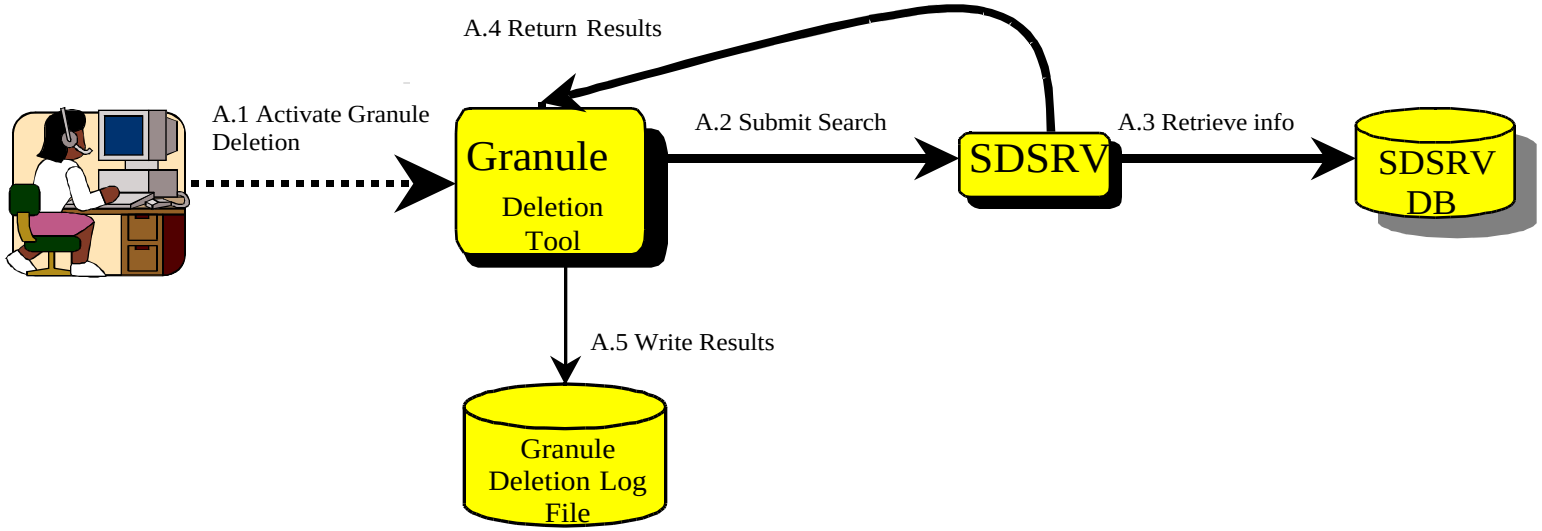

Distributed Object (rpc, CtLib)

HMI (GUI, Xterm, command)

ftp

email (or other as noted)
Double line - Synchronous

email (or other as noted)
Dashed - Asynchronous

Display Candidate Granules for Deletion

End User Interaction

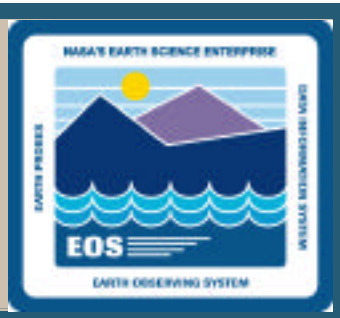

- Physically delete granules and associated granules specified by Short Name and Version ID in the temporal range specified by begindate and enddate.
- Do not prompt for confirmation

- EcDsGranuleDelete

-name <ESDT_ShortName>

-version<ESDT VersionID>

-begindate<date>

-enddate<date>

-log<logfilename>

-physical

-noprompt

End User Interaction

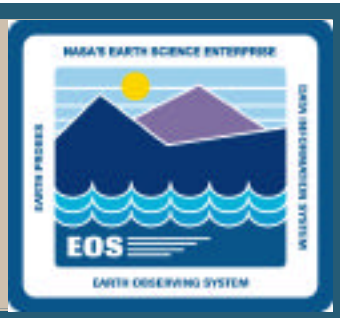

- Delete from the Archive granules specified by Short Name and Version ID that were inserted between insertbegin and insertend
- Prompt for confirmation

- EcDsGranuleDelete

-name <ESDT_ShortName>

-version<ESDT VersionID>

-insertbegin<date>

-insertend<date>

-log<logfilename>

-DFA

End User Interaction

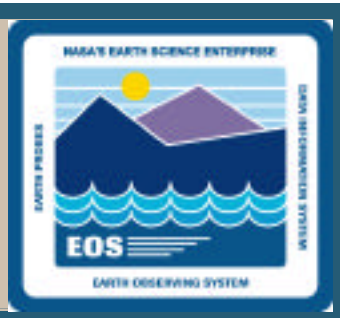

- Delete from the Archive granules specified in the ASCII file that contains (GEOIDs)
  - SC:AST_L1BT.001:1001
- Prompt for confirmation

- EcDsGranuleDelete

-geoidfile<filename>

-log<logfilename>

-DFA

End User Interaction

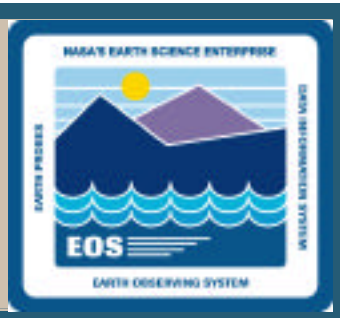

- Delete from the Inventory and the Archive granules and associated granules specified in the ASCII file that contains ShortName, Version ID, and Local Granule ID
  - AST_L1BT 1 tahoe_north_middle
- Do not prompt for confirmation

- EcDsGranuleDelete

-localgranulefile<filename>

-log<logfilename>

-physical

-noprompt

End User Interaction

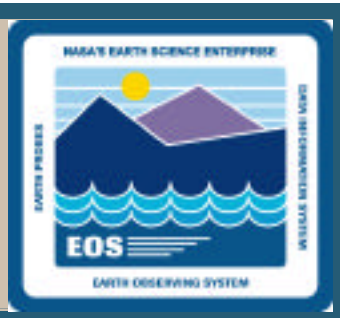

- Display the science granules that are candidates for deletion if the following file was given as input to the Granule Deletion tool

- EcDsGranuleDelete

-localgranulefile<filename>

-display

DAAC Operational Impacts

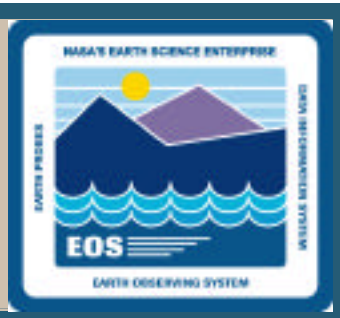

- DAAC Operations must secure this tool using Unix file security
- Must establish procedures to reduce the possibility of deleting granules that are referenced by Processing or still being accessed by the user
- Additional operational procedures will have to be executed to reclaim tapes

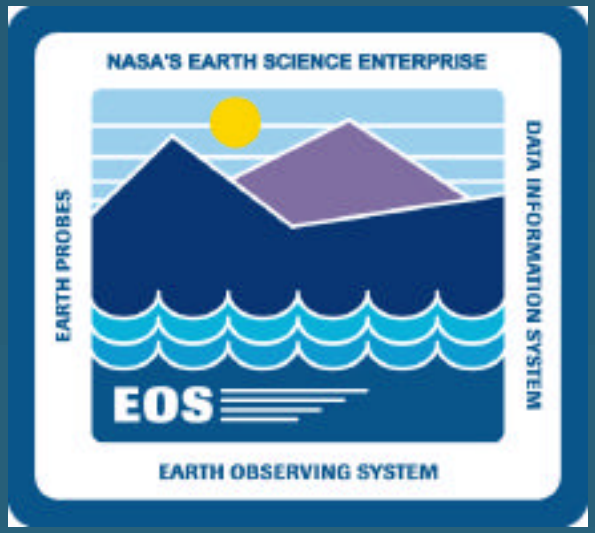

SDSRV Performance Enhancements

John Cockey

Raytheon

Requirements Summary

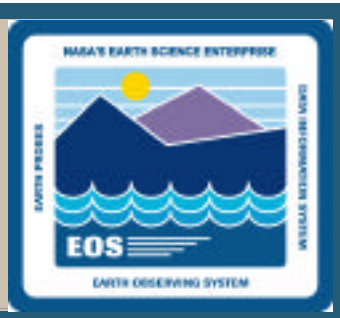

- Driven by capacity requirements outlined in the Drop 6A Release Plan Appendix A - Release 6A Workload Specification
- Key SDSRV related system performance areas:
  - Ingest Criteria
  - Production Criteria (PDPS Performance Briefing)
  - Distribution Criteria
  - Data Access Criteria

Benchmark Methodology

- Decompose Release 6A workload specification to system/subsystem “components” and “key resources”.
- Examine elapsed time of components within
  - SDSRV and V0GW (Search, Integrated Browse)
  - by inference STMGT, DDIST
- [illegible]
  roundtrip and detail level elapsed times
- Execute benchmark runs - Assumptions - single thread, hot cache runs within the EDF environment
- Compare Capacity(current performance) versus Demand (from workload spec.)
- Pinpoint areas where improvement is needed.
- Create 6A performance capabilities to mitigate performance shortfalls.

Key Design Drivers

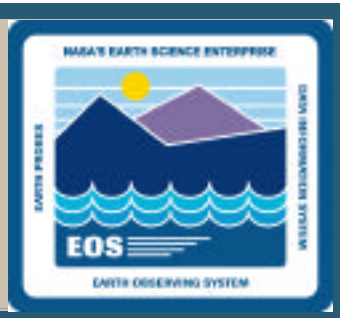

- **Ingest Criteria**
  - Metadata and Data Product Insert
  - Metadata and Data Product Delete
    - Supports Inventory Cleanup, Reprocessing, ESDT specific processing (Landsat 7 DLL)
  - Metadata Update
    - Includes QA metadata update, ESDT specific processing (ASTER Browse)
- **Production Criteria**
  - Metadata and Data Product Insert, Delete, Metadata Search
  - Acquire of Data Products, Inspect of Production Inputs
- **Distribution Criteria - Acquire (including Browse)**
- **Data Access Criteria - Metadata Search, Inspect and Integrated Browse**

Benchmark Status

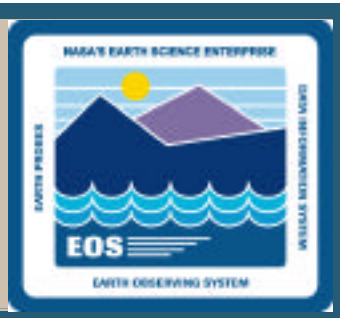

- Completed
  - Decomposition workload specification
  - Instrument code
  - Execute preliminary runs
- Create and prioritize 6A performance capabilities - In process
- Finalize areas for improvement - Nearing Completion
- Measure impacts of improvements
- Iterate - analyze, change, measure

Key Design Drivers
Goals

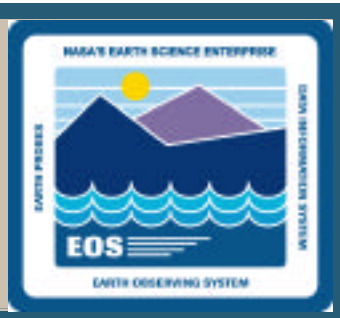

- Overall
  - Reduce Operating System context switching overhead
- Inserts
  - Reduce overall elapsed time by two-thirds+
- Deletes
  - Reduce overall elapsed time by two-thirds+
- Searches
  - Reduce contention to all update activity

Key Design Drivers Target Areas

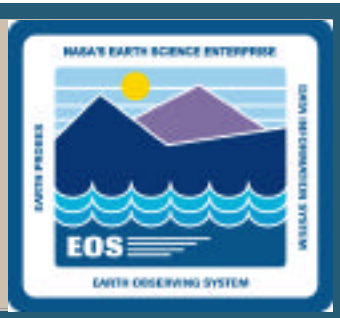

- SDSRV Database Mutex Lock
  - Reduce time spent in lock
- Sybase and Spatial Query Server usage
  - More efficiently use these important resources
- Other SDSRV functions - examples
  - Mass deletes in maintenance window
  - Auto-inspect configurable
  - OS context switching overhead
  - Implement “Read Uncommitted”

SDSRV
Workload Specification Rollup

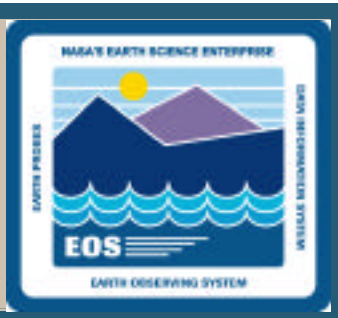

DB Mutex Lock - Demand vs 5B Benchmark Capacity

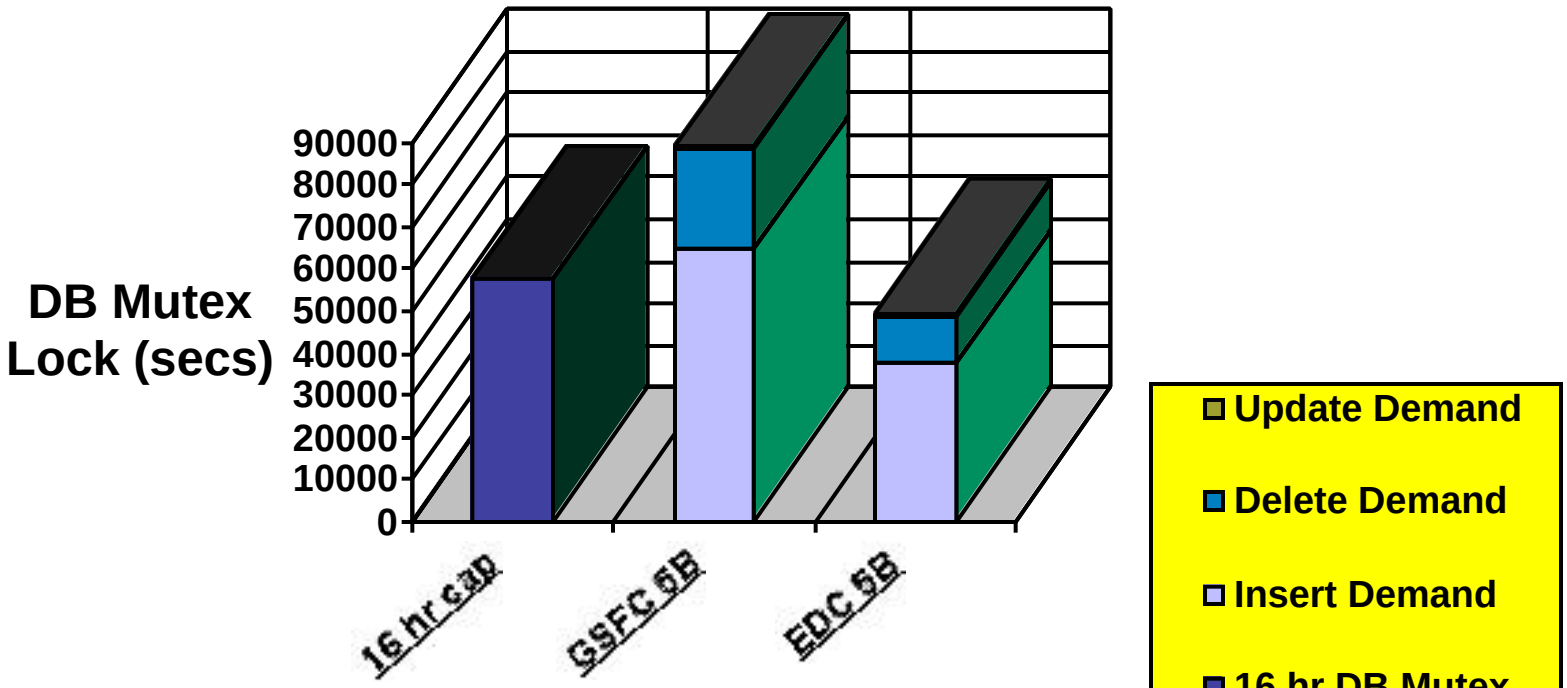

Note - Demand derived from 6A Workload Spec

SDSRV
Workload Specification Rollup

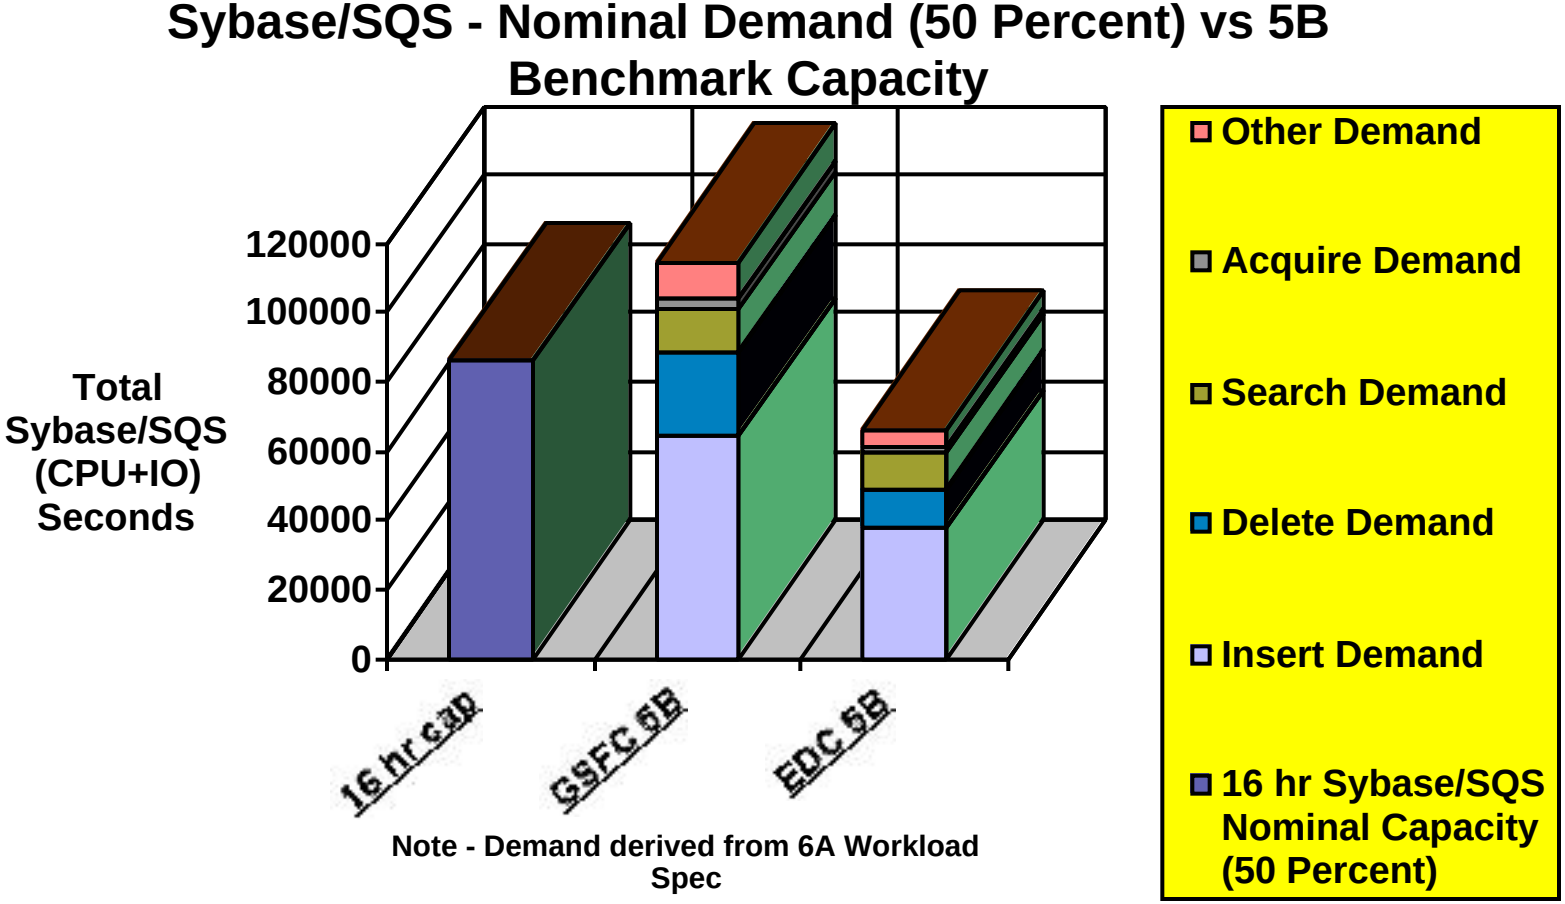

Action Plan
Sample options being considered

1 **Issue: Time spent in DB Mutex during Inserts ~.9 fixed +.05 per PSA seconds* (1.5 for MODAPPS granule; 4.0 for L70RWRS)***

Action: Reduce SQS use - implement 2 phase insert transaction

Estimated Improvement: ~10-15 times improvement

Action: Reduce RPCs to Sybase/SQS - Batch SQL Calls to Sybase

Estimated Improvement : ~ .2 to .4 seconds / granule

2 **Issue : Delete directly contends with Insert - Reduce Impact to "prime time" operations**

Action: logical delete now (timestamp), defer physical delete to scheduled window to reduce DB lock contention/overall resource usage

Estimated Improvement: ~4 times less contention; ~ 3 times better throughput

***Note - see Benchmark Assumptions, slide 3**

Action Plan (continued)

3 **Issue: Current Delete about 3.5 seconds***

Action: Mass delete direct to SQL Server during maintenance window

Estimated Improvement: from 1.7 to about .2 seconds/granule*

4 **Issue: V0GW Inspect overhead high for Searches**

Estimated Improvement: ~.36 seconds/granule returned from a search

5 **Issue: Reduce impact of OS context switching overhead**

Action: Lower impact of malloc()/RWCString overhead

Estimated Improvement: Significant, much better concurrent CPU usage results

***Note - see Benchmark Assumptions, slide 3**

Action Plan (continued)

6 **Issue: Search activity contends with Inserts, Updates for SQL Server locks**

Action: Implement Isolation Level "Read Uncommitted" selectively to reduce DB lock contention

Estimated Improvement: Eliminates DB lock contention for Inserts, Updates; Increases overall throughput by reducing (up to 2) retries from deadlocking

7 **Additional Options**

- **Multiple SDSRV instances**

***Note - see Benchmark Assumptions, slide 3**

SDSRV
5B and 6A Comparison

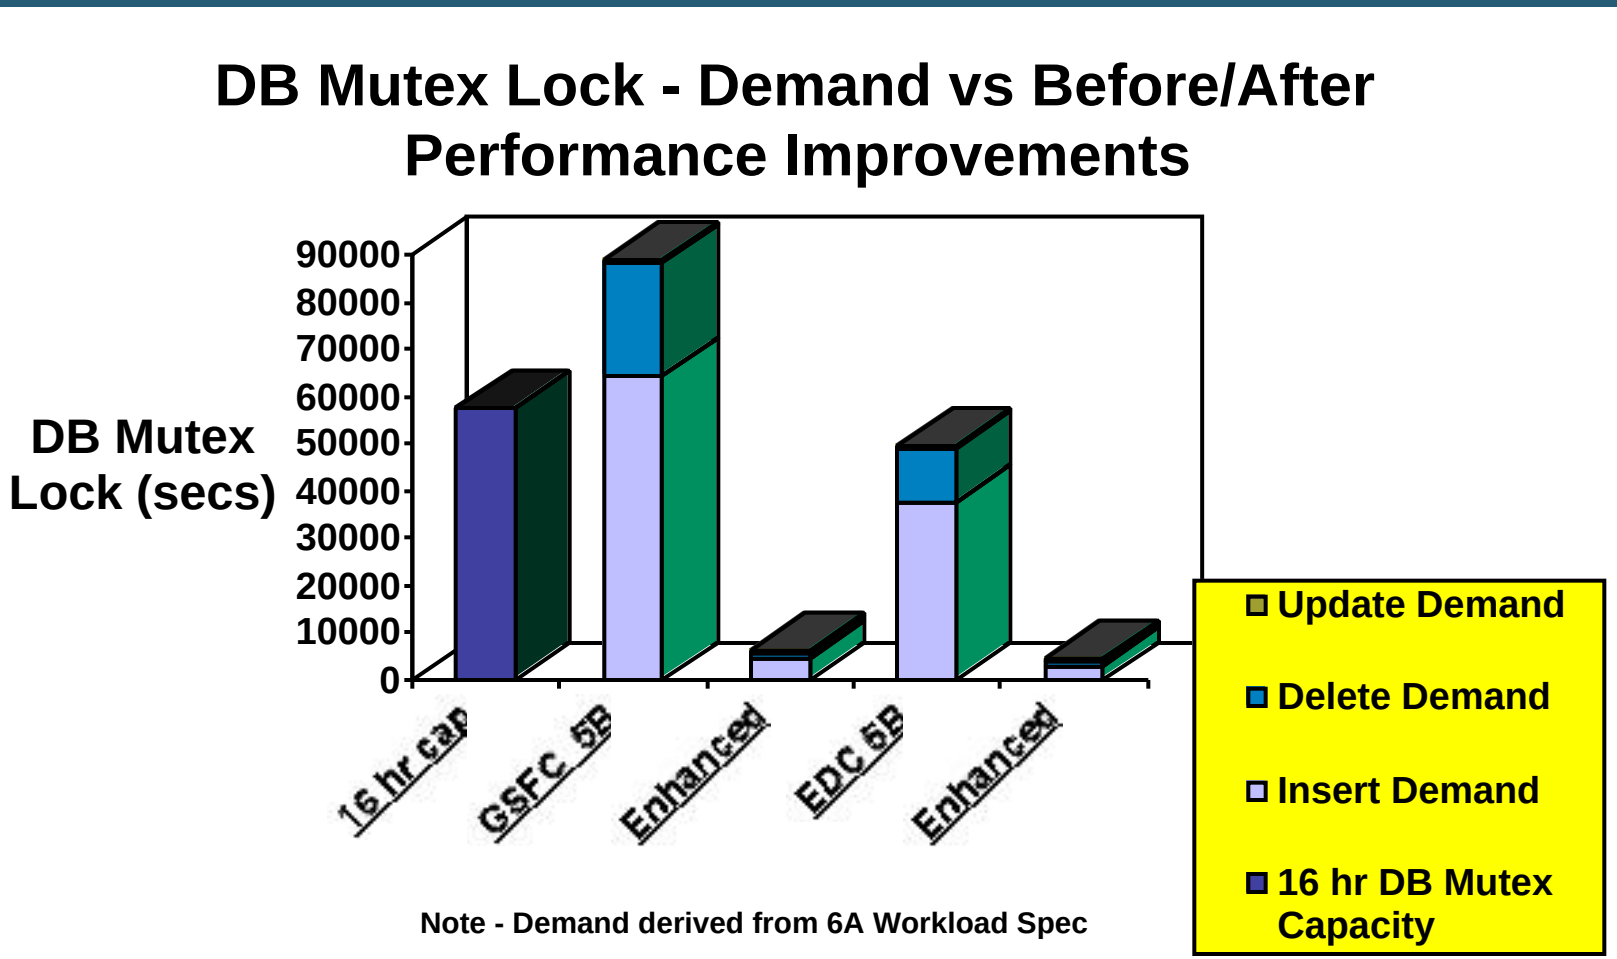

SDSRV
5B and 6A Comparison

Sybase/SQS - Nominal Capacity (50 Percent) vs Demand After Enhancements

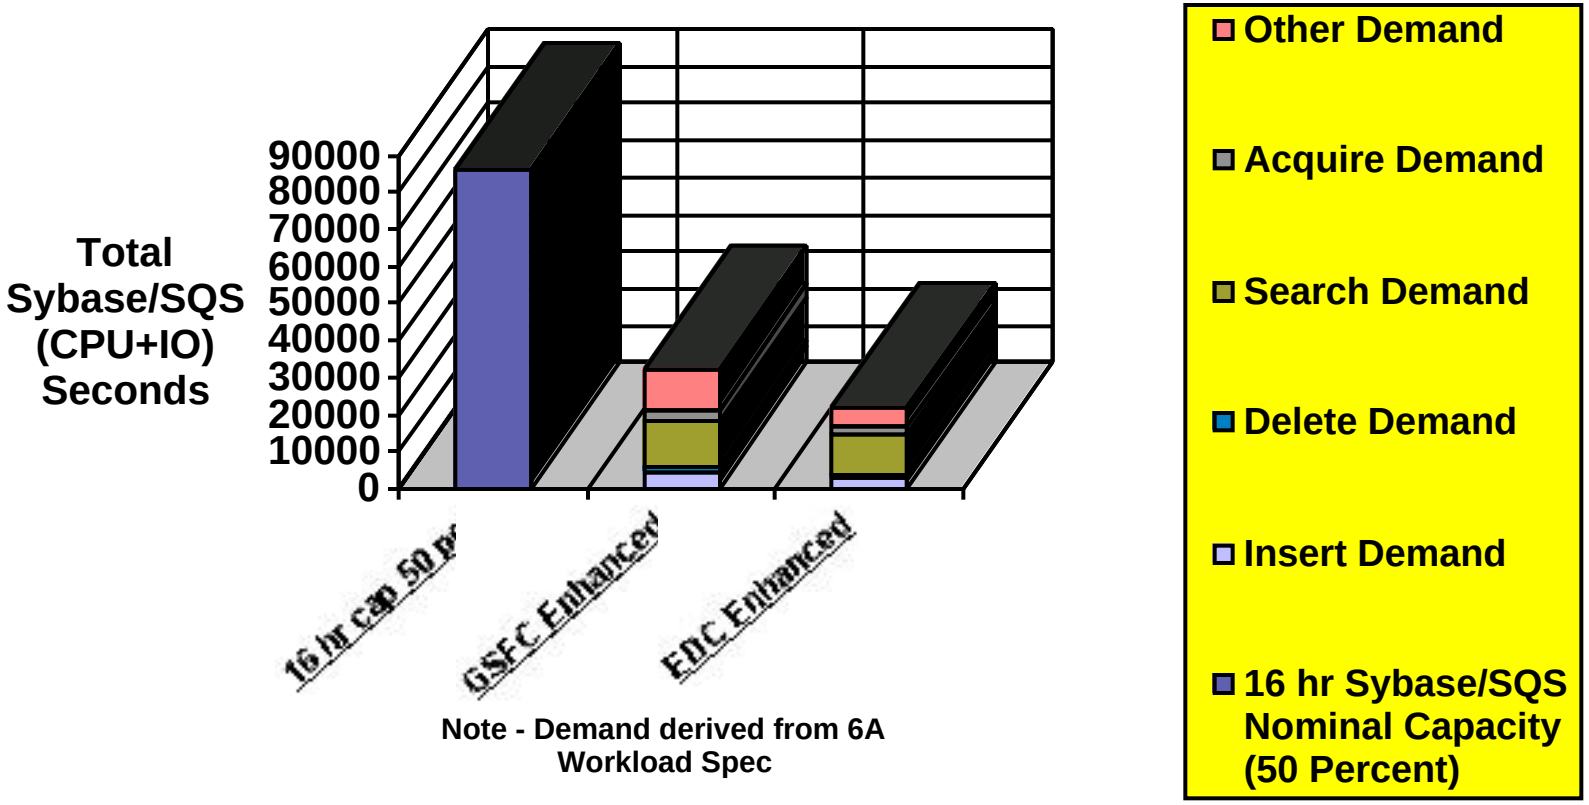

Note - Demand derived from 6A Workload Spec

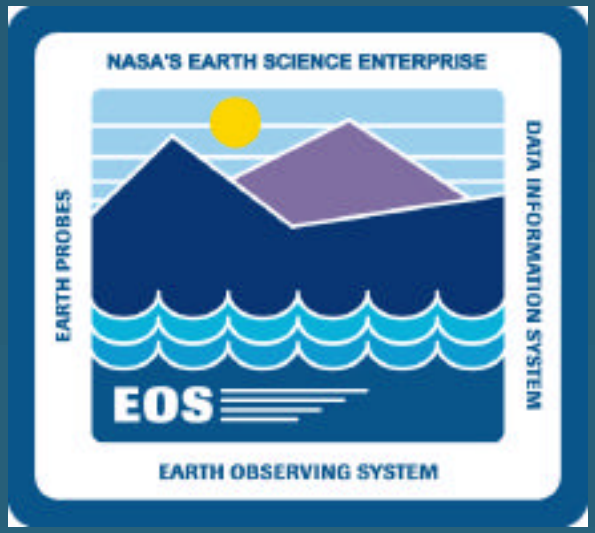

PDPS Performance Improvements

Desryn Duncan

Raytheon

Requirements Summary

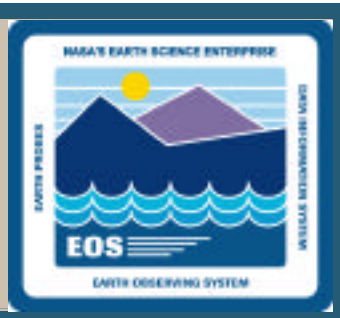

- Driven by capacity requirements outlined in the Drop 6A Release Plan Appendix A - Release 6A Workload Specification

- Key PDPS related system performance areas:
  - Planning:
    - Generate Data Processing Requests (DPR) from entered Production Requests
    - Process Subscription Notifications & Release DPR
  - Processing:
    - Allocate & De-allocate Processing Resources for DPR
    - Submit Staging & Destaging Requests to DSS

Key Design Drivers
Methodology (PLS)

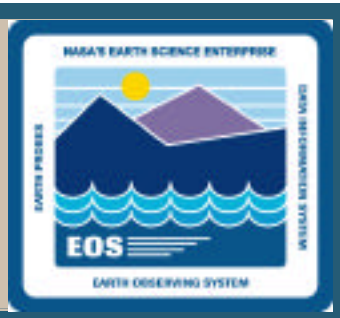

Objective: Produce Daily Plan in 8 Hours or Less

Executed Benchmarks using Instrumented Code

- **Divided code into logical bins**
- **Measured cpu and wall-clock time for each bin**
- **Logged measurements to file**
- **Parsed and summed up results using scripts**

Compared Capacity with Targets

Determined Key Areas for Improvement

Identified Technical Resolution

Key Design Drivers
Objectives/Current Capacity (PLS)

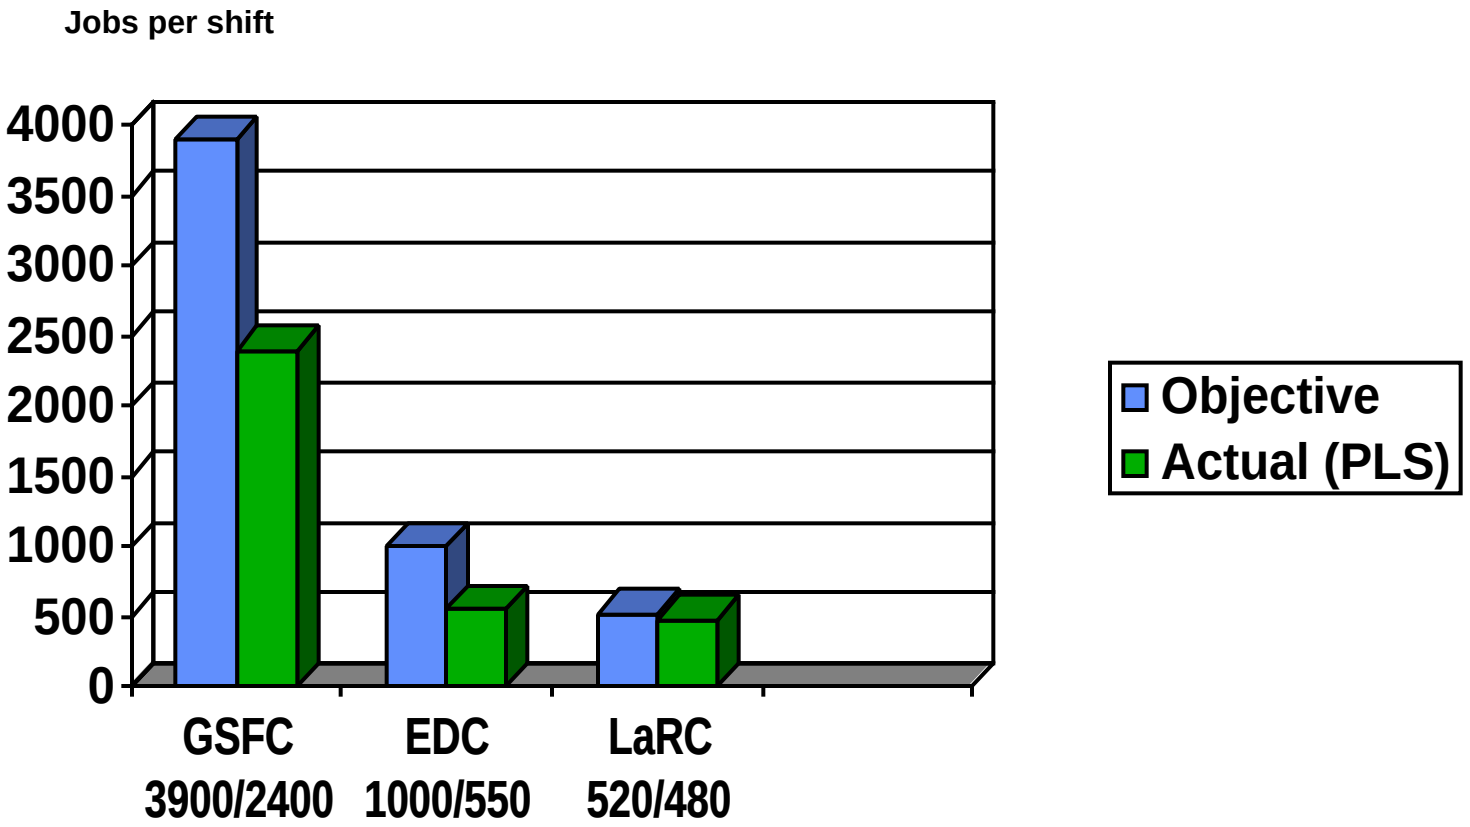

Key Design Drivers Action Plan (PLS)

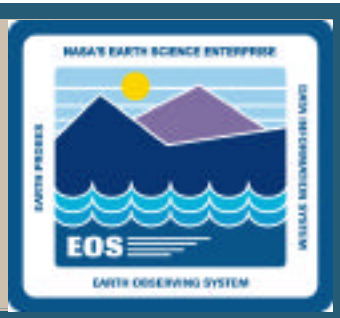

- **Issue: Populating DB with DPR inputs and outputs takes a long time (80-90% of Planning Time spent here!)**

  Action: Use stored procedures, streamline code, reduce repetitive object construction

  Estimated Improvement: 300% for MODIS

- **Issue: Inefficient sorting of SDSRV search results**

  Action: Improve algorithm

  Estimated Improvement: 15%

PLS
5B and 6A Comparison

Demand vs Est. 6A Enhanced Capacity

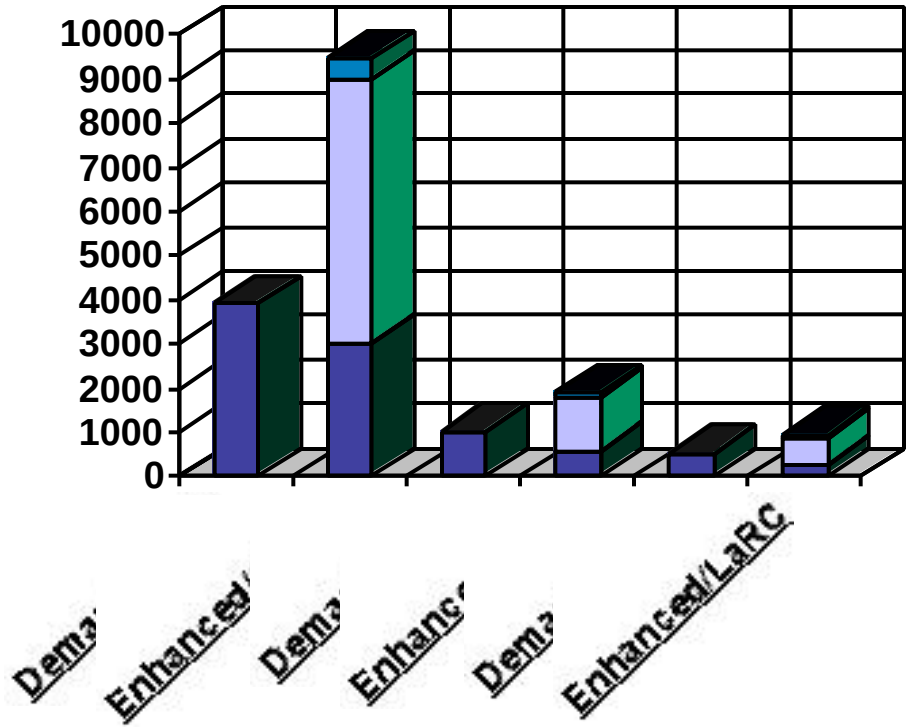

- Improve Sort Alg.
- Stored Procedures
- Demand

Note - Demand derived from 6A Workload Spec

Key Design Drivers
Methodology (DPS)

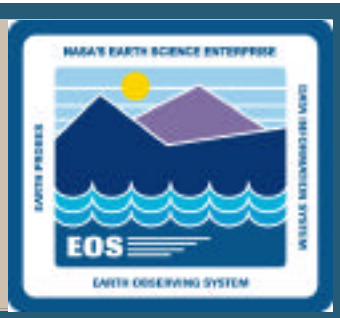

Objective: Throughput of DPS Components Equal or Better Daily Processing Workload

Developed DPS Simulator To Benchmark Current Capacity and Improvements

- **PDPS database loaded with required PLS parameters**
- **DSS services stubbed out. Files retrieved from a directory**
- **Simple PGE script verifies accuracy of PCF file and generates output files specified in PCF**
- **Test measures throughput and loading of resources**

Identified Promising Improvements

Determined Technical Approaches

Key Design Drivers
Objectives/Current Capacity (DPS)

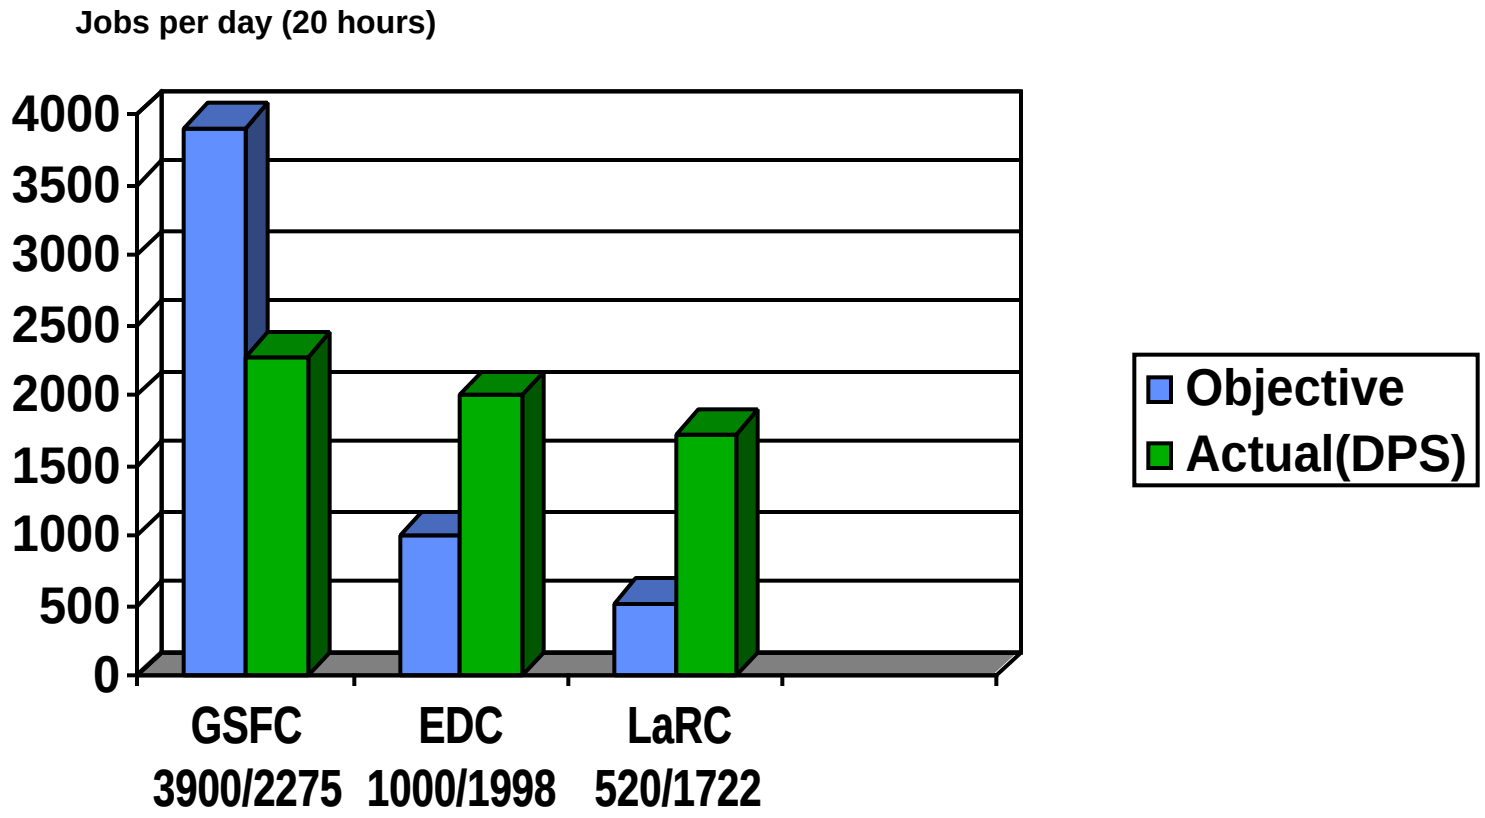

Key Design Drivers Action Plan (DPS)

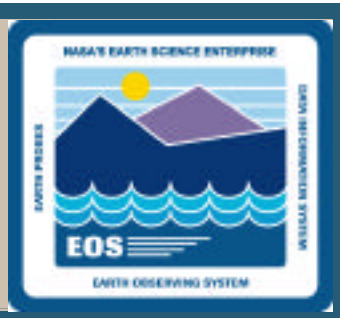

- **Issue: Resource Manager Overhead & Contention**
  Action: Replace Lock Manager with stored procedures
  Estimated Improvement: 35%

- **Issue: CPU Capacity Is Bottleneck**
  Action: Increase CPU capacity to 4 x 336 MHz
  Estimated Improvement: 30%

- **Issue: Autosys Event Handler Degrades As Jobs Increase**
  Action: Combine jobs in Autosys job box
  Estimated Improvement: 40%

Key Design Drivers Action Plan (DPS)

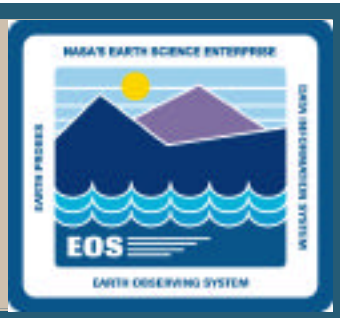

- **Additional Options Available**

  1. **Add Autosys instance**

  2. **Add Queueing Server boxes**

DPS
5B and 6A Comparison

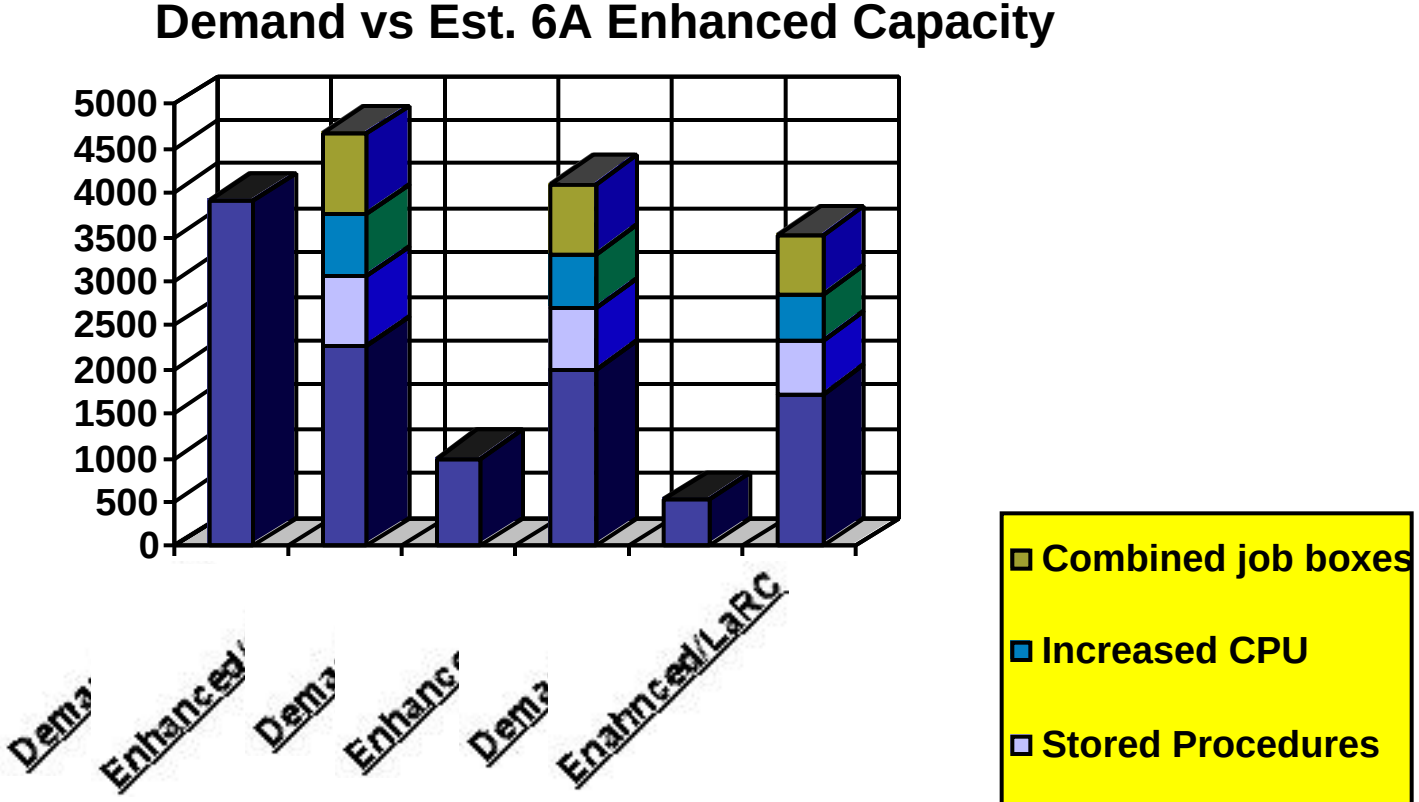

Note - Demand derived from 6A Workload Spec

Other PDPS Performance Enhancements

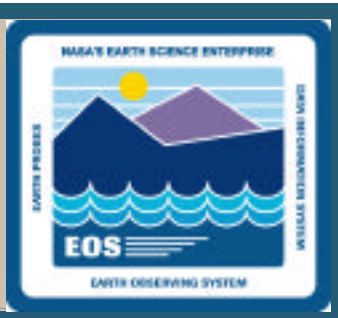

- **Subscription Notification**
  Issue: Checking whether a DPR is ready to run takes a long time

- **Deletion of DPRs**
  Issue: The deletion process for DPRs takes a long time

Still evaluating both issues, but early results indicate gains will be made by using more stored procedures and improving the relevant algorithms.